

Bayesian approach to LR assessment in case of rare type match: careful derivation and limits

Giulia Cereda,^{1*,2}

¹University of Lausanne, Faculty of Law, Criminal Justice and Public Administration

²Leiden University, Mathematical Institute

*To whom correspondence should be addressed; E-mail: giulia.cereda@unil.ch.

The likelihood ratio (LR) is largely used to evaluate the relative weight of forensic data regarding two hypotheses and for its assessment Bayesian methods are widespread in the forensic field. However, the Bayesian ‘recipe’ for the LR presented in most of literature consists in plugging-in Bayesian estimates of the involved nuisance parameters into a frequentist-defined LR: frequentist and Bayesian methods are thus mixed, giving rise to solutions obtained by hybrid reasoning. This paper provides the derivation of a proper Bayesian approach to assess LR for the ‘rare type match problem’, the situation in which the expert wants to evaluate a match between the profile of a suspect and that of a trace from the crime scene, and this profile has never been observed before in the database of reference. Bayesian LR assessment using the two most popular Bayesian models (beta-binomial and Dirichlet-multinomial) is discussed and compared to corresponding plug-in versions.

Key words: Bayesian plug-in, evidence evaluation, rare type match, Y chromosome STR, beta binomial model, Dirichlet-multinomial model, hierarchical bayesian model.

1 Introduction

One of the main challenges of forensic science is that of properly evaluating the match between the characteristics of a crime stain (for instance a DNA profile) and the corresponding characteristics of some material from a known source (for instance from a suspect). Typically, a couple of mutually exclusive hypotheses is defined, of the kind of ‘the crime stain came from the suspect’ (h_p) and ‘the crime stain came from an unknown donor’ (h_d).

The forensic expert is given some data D which can typically be split into *evidence*, data directly related to the crime, and *background*, additional data not directly related to the crime and only pertaining some nuisance parameter θ involved in the assessment of the likelihood ratio. Evidence and background data will be modeled through random variables E and B respectively. In particular, we are interested in the situation in which the forensic expert is asked to evaluate the match between a DNA profile of a suspect and the DNA profile of a stain found at the crime scene. It is intuitive to understand that (one of) the nuisance parameter(s) involved in this evaluation is the proportion of people with the same profile among the possible perpetrators: the more this profile is rare the more the suspect is in trouble. This parameter is unknown and thus the expert is given (or asks for) a database containing a list of DNA profiles from a sample from the population of possible perpetrators. The main difference between the frequentist and the Bayesian methodology is that the first considers the nuisance parameter θ and the correct hypothesis h as fixed (without distribution) unknown quantities, while the second models the expert's uncertainty about their value through random variables, whose prior distributions reflect prior expert's beliefs.

The largely accepted method for evaluating the data in order to discriminate between the two hypotheses of interest, is the calculation of the *Bayes factor* (BF), regularly called in forensic context *likelihood ratio* (LR) and defined as the ratio of the probability of observing the data under the two competing hypotheses:

$$\text{LR} = \frac{\Pr(E = e, B = b \mid H = h_p)}{\Pr(E = e, B = b \mid H = h_d)}.$$

In the Bayesian framework (the one of interest for this paper) \Pr is the joint distribution of all the random variables in the model (E , B , H , and Θ).

On the other hand, frequentists, which consider θ and h as fixed quantities, use a different probability (here denoted as $\mathcal{P}r$) which can be expressed in terms of the Bayesian \Pr , in the following way: $\mathcal{P}r(\cdot) := \mathcal{P}r_h^\theta(\cdot) = \Pr(\cdot \mid \Theta = \theta, H = h)$. Thus, the frequentist likelihood ratio (denoted as $\mathcal{L}\mathcal{R}$) is defined as

$$\mathcal{L}\mathcal{R} = \frac{\mathcal{P}r_{h_p}^\theta(E = e, B = b)}{\mathcal{P}r_{h_d}^\theta(E = e, B = b)} = \frac{\Pr(E = e, B = b \mid \Theta = \theta, H = h_p)}{\Pr(E = e, B = b \mid \Theta = \theta, H = h_d)}.$$

Depending on the preferences of the expert, frequentist or Bayesian likelihood ratios can be used for the evaluation of forensic data. Once a choice has been made, it is important to be consistent with it, while literature often mixes up the two. To our knowledge, this paper and (Cereda, 2015b) constitute the only forensic literature that discuss the differences between the two approaches. (Cereda, 2015b) is concerned with the theoretical foundations of frequentist solutions, while this paper provides a simple and careful derivation of the proper Bayesian LR, for the rare type match problem (described in Section 2): the situation

in which the DNA profile of the crime stain and that of the suspect match but they are not among the DNA profiles observed in the reference database. In Section 3 we will discuss the fact that influential Bayesian forensic literature (Weir, 1996; Aitken and Taroni, 2004; Taroni et al., 2010, 2014; Sjerps et al., 2015) seems to suggest the use of frequentist defined likelihood ratio (\mathcal{LR}) and use Bayesian methodologies only inasmuch they provide a Bayesian estimate of θ to be plugged into \mathcal{LR} . Others (Curran, 2005; van der Hout and Alberink, 2015), treat the likelihood ratio as function of θ and provide its posterior distribution with respect to the posterior distribution of θ given the data. However, one of the main points of discussion is that there is no need of such hybrid derivations since the proper Bayesian LR is often very easy to obtain: this paper shows how this should be done, taking advantage of a very useful Lemma, presented in Section 4. However, for this method to be advisable, the Bayesian prior should be chosen in a sensible way, reflecting the expert's opinion, and not by mathematical convenience as often happens.

The two most common Bayesian models (beta-binomial and Dirichlet-multinomial) are discussed in Sections 5 and 6. They are general enough to be applied to different kinds of forensic evidence evaluation, but will be here applied to DNA profiles obtained using the Y-STR marker system, with the double aim of exploring the performance of the conventional Bayesian prior choices for the rare type match case for non autosomal DNA, and of showing how a full Bayesian LR is to be defined and calculated. Sensitivity analysis and comparison with proposed hybrid plug-in solutions are carried out. We are not entirely satisfied with the performance of classical models for the rare haplotype problem, which we believe would need different kinds of prior, more realistic and tuneable, such as those proposed in Cereda (2015c).

1.1 Notation

Throughout the paper the following notation is chosen: random variables and their values are denoted, respectively, with uppercase and lowercase characters: x is a specific realisation of X . Random vectors and their values are denoted, respectively, by uppercase and lowercase bold characters: \mathbf{p} is a realisation of the random vector \mathbf{P} . Bayesian probability is denoted with $\Pr(\cdot)$, while density of a continuous random variable X is denoted by $p(x)$. For a discrete random variable Y , the continuous notation $p(y)$ and the discrete one $\Pr(Y = y)$ will be alternately used. Frequentist probability will be denoted as \mathcal{P} .

2 The rare type match problem

The DNA sequence of an individual is a very long sequence of letters (each corresponding to 4 nucleotides) which code the genetic instruction necessary for the life of the individual. The entire sequence is unique to each individual (with the only exception of homozygous

twins, which share the same sequence), but a Y-STR DNA profile (also called *haplotype*) usually consists on a short list of integers (typically 7 to 23) that describes only certain characteristics of the DNA sequence of the individual on the Y chromosome (Gill et al., 2001). Moreover, the Y-STR profile is shared between men in the same patrilineal lineage. For these reasons there is a positive probability that two different persons share the same profile. This is why we need to weight how probable is the observed match under the hypothesis that the suspect left the stain against how likely is the match under the hypothesis that someone else left the stain. Clearly, assuming that a match is always detected correctly (no false positives), the first probability is 1, and the second depends on the proportion θ of that profile among the population of potential perpetrator. Moreover, we are given a list of profiles from a sample of individuals belonging to the population of possible perpetrators to assess this frequency. Problems arise when the observed frequency of this characteristic is 0, the so-called ‘rare type match problem’. This problem is particularly significant in case a new kind of forensic evidence for which the available database size is still limited is involved. This is the case, for instance, when using DIP-STR markers (Cereda et al., 2014)). The same happens when Y-chromosome (or mitochondrial) DNA profiles are used, since the set of possible haplotypes is extremely large, and the coverage of available databases often limited. The case of Y-STR DNA will thus be retained here as an extreme but in practice common and important way in which the problem of assessing the evidential value of rare type match can arise. This is a very appropriate and paradigmatic example, since literature provides examples of different approaches to evaluate the evidential value of rare Y-STR profile match (Roewer et al., 2000; Andersen et al., 2013, e.g.), even though, in our opinion, a proper Bayesian derivation for the LR in rare type match case hasn’t been proposed yet. This problem is so substantial that it has been defined “the fundamental problem of forensic mathematics” by Brenner (2010). We will now review some of the methods proposed by literature. Most of them address the problem of assessing the frequency of a type with zero occurrence, sometimes under the name of ‘zero numerator problem’ (e.g. Winkler et al., 2002). Notice that this is related, but not equivalent, to the problem of assessing the likelihood ratio in case of a rare type match. The *empirical frequency estimator*, also called *naïve estimator*, that uses the frequency of the characteristic in the database, puts unit probability mass on the set of already observed characteristics, and it is thus unprepared for the observation of a new type. A solution could be the *add-constant* estimators (in particular the well known *add-one* estimator, due to Laplace (1814), and the *add-half* estimator of Krichevsky and Trofimov (1981)), which add a constant to the count of each type, included the unseen ones. However, this method requires to know the number of possible unseen types, and does not perform well when this number is large compared to the sample size (see Gale and Church (1994) for additional

discussion). Moreover, [Louis \(1981\)](#) proposes the so-called ‘rule of three’, that states that if n is the size of the database, $3/n$ is a good approximation of the 95% upper bound for the frequency. This is also proposed in a Bayesian framework, by [Jovanovic and Levy \(1997\)](#); [Winkler et al. \(2002\)](#); [Chen and McGee \(2008\)](#). Alternatively, [Good \(1953\)](#), based on an intuition on A.M. Turing, proposed the nonparametric *Good Turing estimator* for the total unobserved probability mass, based on the proportion of singleton observations in the sample. An extension of this estimator is applied to the LR assessment in the rare type match in [Cereda \(2015b\)](#). For a comparison between *add one* and *Good-Turing* estimator, see [Orlitsky et al. \(2003\)](#). As pointed out in [Anevski et al. \(2013\)](#), the *naive estimator*, and the *Good Turing estimator* are in some sense complementary: the first gives a good estimate for the observed types, and the second for the probability mass of the unobserved ones. More recently, [Orlitsky et al. \(2004\)](#) have introduced the *high profile estimator*, which extends the tail of the *naive estimator* to the region of unobserved types. [Anevski et al. \(2013\)](#) improved this estimator and provided the consistency proof. Papers that address the rare Y-STR haplotype problem in forensic context are for instance [Egeland and Salas \(2008\)](#), [Brenner \(2010\)](#), and [Cereda \(2015a,c\)](#). Moreover, the Discrete Laplace method presented in [Andersen et al. \(2013\)](#), even though not specifically designed for the rare type match can be successfully applied to that extent ([Cereda, 2015b](#)). Bayesian nonparametric estimators for the probability of observing a new type have been proposed by [Tiwari and Tripathi \(e.g. 1989\)](#); [Lijoi et al. \(e.g. 2007\)](#); [Favaro et al. \(e.g. 2009\)](#). However, for the likelihood ratio assessment it is required not only the probability of observing a new species but also the probability of observing this same species twice (according to the defense the crime stain profile and the suspect profile are two independent observations). [Cereda \(2015c\)](#) is the first paper that addresses the problem of LR assessment in the rare haplotype case using Bayesian nonparametric models.

3 The full Bayesian approach to LR

The likelihood ratio assessment often involves some unknown nuisance parameters, denoted as θ . In our case, it is the proportion of individuals in the relevant population with Y-STR profile corresponding to that of the matching trace, or the entire vector containing the frequencies of all the profiles. The parameter of interest, h , is the unknown true hypothesis. Available data is made of evidence (E) directly related to the crime, and which helps in discriminate h , and additional background data (B) not directly related to the crime and only pertaining to the nuisance parameter θ . This is partially different from the ‘background information’ I as defined in [Aitken and Taroni \(2004\)](#); [Taroni et al. \(2014\)](#), but we can say that often background data can be thought of as part of the background information.

The difference between Bayesian and frequentist methods consists in how they treat the parameters θ and h . A Bayesian models the uncertainty about their value by random variables Θ and H , which are given prior distributions $p(\theta)$ and $p(h)$. Frequentists consider them as fixed (i.e., without distribution) unknown quantities. The reader is invited to notice the difference between θ and h : one is the parameter which we ‘test’ through the likelihood ratio (h), the other (θ) is a nuisance parameter involved in its calculation. Some assumptions about the conditional independence probability for the model can be made, valid both for the frequentist and for the Bayesian approach.

- a. The distribution of B given h and θ , only depends on θ .
- b. B is independent of E , given θ and h .

In our DNA example, condition **a** corresponds to ask that the sampling mechanism to obtain the database of reference is independent of which hypothesis is correct. This is true if, as it often happens, the database is collected before the crime. Condition **b** holds if the suspect has been found on a ground of different evidence that has nothing to do with DNA. In what follows we are going to use Bayesian network notation to specify the conditional independence relations of the proposed models. We expect the reader to be familiar with such a representation.

3.1 Bayesian point of view.




Figure 1: Bayesian network representing the dependency relations between E (evidence of the case) B (background data in the form of a database) Θ (population parameter) and H (hypotheses of interest).

Bayesians deal with the uncertainty over the parameters θ and h by considering their values as realisations of, respectively, random variables Θ and H . A full Bayesian model is defined when the prior joint probability distribution for all the random variables of the model (here E , B , H and Θ) is given. This full Bayesian model can be thus represented by the Bayesian network of Figure 1, which is in turn equivalent to the following three conditions:

- Bayesian a.** B is conditionally independent of H given Θ .

Bayesian b. B is conditionally independent of E given Θ and H .

Bayesian c. Θ is unconditionally independent of H .

Notice that they are the Bayesian reformulation of conditions **a.** and **b.** mentioned above, with an additional condition (**Bayesian c.**) which corresponds to assuming that the Bayesian probability makes Θ and H independent, and is guaranteed for instance if prior beliefs on θ and on h are assessed by people with different responsibilities and tasks: a judge for h and a DNA expert (or a statistician) for θ . However, notice that by definition the LR is independent of the prior belief over h .

The structure of the Bayesian network (or, equivalently, the three conditions above) allows to factorise the joint prior as $p(\theta, h, b, e) = p(\theta)p(h)p(b|\theta)p(e|\theta, h)$. The Bayesian probability \Pr underlying the model is defined accordingly. As all Bayesian probabilities it is an expression of the subjective belief of the experts. This is achieved by choosing the prior distribution for θ and h which reflects expert's beliefs. The distribution of all other variables given θ and h is defined by the model, and need no subjective assessment.

The Bayesian likelihood ratio can be derived in the following way:

$$\begin{aligned} \text{LR} &= \frac{\Pr(E = e, B = b | H = h_p)}{\Pr(E = e, B = e | H = h_d)} = \frac{\Pr(E = e | B = b, H = h_p)}{\Pr(E = e | B = b, H = h_d)} = \frac{\int p(e|b, h_p, \theta) p(\theta|b, h_p) d\theta}{\int p(e|b, h_d, \theta) p(\theta|b, h_d) d\theta} \\ &= \frac{\int \theta p(\theta|b) d\theta}{\int \theta^2 p(\theta|b) d\theta} = \frac{\mathbb{E}(\Theta | B = b)}{\mathbb{E}(\Theta^2 | B = b)}. \end{aligned}$$

where some simplifications have been carried out because of conditions **a.**, **b.**, and **c.** Moreover, the first equality is due to the fact that it is implied by the network structure that B is also unconditionally independent of H , and the second to last equality is due to the fact that, given condition **a** it holds that $p(\theta|b, h) = p(\theta|b)$, $\forall h$.

3.2 Frequentist point of view.

As already mentioned, frequentists consider h and θ as fixed quantities, whose unknown values correspond to, respectively, the true value of θ and the correct hypothesis. The frequentist model can be thus seen as a special case of the Bayesian model described in Section 3.1, where Θ and H are given degenerate priors on θ and h , respectively. Alternatively, one can express the frequentist probability \mathcal{Pr} in terms of the Bayesian \Pr in the following way: $\mathcal{Pr}(\cdot) := \mathcal{Pr}_h^\theta(\cdot) = \Pr(\cdot | H = h, \Theta = \theta)$. If the Bayesian \Pr was subjective, the frequentist \mathcal{Pr} is a measure which is universally defined by Nature.

Regarding h , according to prosecution its true value is h_p , while according to defence it is h_d . So one can think of two different frequentist probabilities: one for the prosecution ($\mathcal{Pr}_{h_p}^\theta$) and one for the defence ($\mathcal{Pr}_{h_d}^\theta$).

From a frequentist point of view, conditions **a** and **b** correspond to ask that:

Frequentist a. $\Pr_{h_p}^\theta(B = b) = \Pr_{h_d}^\theta(B = b)$, for all θ and b .

Frequentist b. $\Pr_h^\theta(E = e | B = b) = \Pr_h^\theta(E = e)$, for all θ, h, e , and b .

Obviously, **Bayesian c** becomes irrelevant in the frequentist framework.

The frequentist \mathcal{LR} can be derived as:

$$\mathcal{LR} = \frac{\Pr_{h_p}^\theta(E = e, B = b)}{\Pr_{h_d}^\theta(E = e, B = b)} = \frac{\Pr_{h_p}^\theta(E = e | B = b) \Pr_{h_p}^\theta(B = b)}{\Pr_{h_d}^\theta(E = e | B = b) \Pr_{h_d}^\theta(B = b)} = \frac{\Pr_{h_p}^\theta(E = e)}{\Pr_{h_d}^\theta(E = e)} \quad (1)$$

where the last equality is due to conditions **Frequentist a**, and **Frequentist b**.

Stated otherwise, frequentists look at a value for $\mathcal{LR}|\theta$ (read “LR given θ ”), where the value θ is fixed and has to be estimated through data.

Through observations, frequentists attempt to get close to the true \mathcal{LR} by choosing some estimator $\widehat{\mathcal{LR}}$. One possibility is to estimate θ with a particular $\hat{\theta}$. This leads to the so-called *plug-in estimation* $\widehat{\mathcal{LR}} = \mathcal{LR}(\hat{\theta})$ of the \mathcal{LR} . However, that’s not the only option (Cereda, 2015b).

By looking at (1) the reader will realise that, if the frequentist approach is chosen, and under conditions **a** and **b**, one would get to the same result by evaluating only E or both E and B . This means that part of the information, namely B , is not useful to discriminate between the two hypotheses of interest (however, it usually plays an important role to obtain the estimate $\hat{\theta}$ to be plugged into the \mathcal{LR}). The same does not hold in the Bayesian context.

3.3 The Bayesian plug-in LR and the proper Bayesian LR

It is now time to discuss the fact that important forensic literature (e.g. Evett and Weir, 1998; Balding, 2005; Lucy, 2005) considers the likelihood ratio as ‘a measure of the probative value of the evidence regarding the two hypotheses’ h_p and h_d . According to this, it indicates the extent to which E (and only E) is in favour of one hypothesis over the other. This is, in my opinion, the first important problem, since all data at disposal (namely E and B) should be evaluated. Even though this is irrelevant in the frequentist framework (see (1)), in the Bayesian framework for this definition to be appropriate we need to replace the probability \Pr with the posterior probability $\Pr^*(\cdot) = \Pr(\cdot | B = b)$. Indeed, it holds that

$$\mathcal{LR} = \frac{\Pr(E = e, B = b | H = h_p)}{\Pr(E = e, B = b | H = h_d)} = \frac{\Pr(E = e | B = b, H = h_p)}{\Pr(E = e | B = b, H = h_d)} = \frac{\Pr^*(E = e | H = h_p)}{\Pr^*(E = e | H = h_d)}.$$

It is as if we have separated the evaluation process in two steps: first we observe $B = b$, and update the probability $\Pr(\cdot)$ to the posterior $\Pr^*(\cdot) = \Pr(\cdot | B = b)$, and then we define the likelihood ratio as the ratio of the probabilities (\Pr^*) of observing (only) the evidence

E , under the two alternative hypotheses.¹ This point is generally mistaken in literature and, as a consequence, the problem is split into two phases: first, a Bayesian estimate of θ using B , in the form of a posterior expectation is obtained, and then this estimate is plugged into a frequentist defined \mathcal{LR} . It is as if, instead of using a combined model such as that in Figure 1, they use two separate models as those in Figure 2: the left one is used to update the prior over the parameter. The second one is used to derive the likelihood ratio (with θ considered as a fixed quantity). In the end, they replace θ with the posterior expectation of Θ (now modelled through a random) given B . This method will be referred to in the paper as the ‘Bayesian plug-in method’, since it is wrongly considered Bayesian, but it actually plugs in Bayes estimates into likelihood ratio defined in a frequentist way.




Figure 2: The two phase approach corresponding to Bayesian plug-in.

The correct Bayesian approach would be either to evaluate both E and B simultaneously, using the network of Figure 1, or in two steps: after the observation of B , we can update the model to the one represented in Figure 3, and use this for the evaluation of E .




Figure 3: Updated Bayesian network after the observation of B .

3.4 State of the art for DNA match evaluation

In case of a DNA match, we can use the Bayesian network of Figure 4, which is equivalent to the network in Figure 1 with the only difference that here node E is split into two separated nodes, E_s and E_c representing the suspect’s and the crime stain’s profile, respectively. We

¹ Often, in literature (Taroni et al., 2014, e.g.), it is explicitly stated that I , the so-called background information, is omitted in the notation. We then agree with this choice provided that B is part of I .

denote with θ the unknown vector made of the population proportions of the different DNA profiles in the population, modelled through the random variable Θ . Here we assume that we know the whole list of different DNA types present in the population of possible perpetrators, later we will consider the situation in which we don't. With Θ_{e_s} we will denote the population frequency of the suspect's (and crime stain's) profile.




Figure 4: Bayesian network for the DNA example.

According to the prosecution, the suspect left the stain. This implies that $\Pr(E_c = e_s | \Theta = \theta, E_s = e_s, H = h_p) = 1$, under the assumption that each true match is correctly reported. According to the defence, another person from the population left the stain, hence the probability of it being exactly of type e_c is equal to the population proportion of that profile: $\Pr(E_c = e_s | \Theta = \theta, E_s = e_s, H = h_d) = \theta_{e_s}$. Moreover, it holds that $p(b|e_s, \theta) \Pr(e_s|\theta)p(\theta)$ is proportional to $p(\theta|e_s, b)$. The correct Bayesian procedure would lead to:

$$\begin{aligned} \text{LR} &= \frac{\Pr(E = e, B = b | H = h_p)}{\Pr(E = e, B = b | H = h_d)} = \frac{\int \Pr(E_c = e_s | H = h_p, \Theta = \theta, E_s = e_s) p(e_s|\theta) p(b|\theta, e_s) p(\theta) d\theta}{\int \Pr(E_c = e_s | H = h_d, \Theta = \theta, E_s = e_s) p(e_s|\theta) p(b|\theta, e_s) p(\theta) d\theta} \\ &= \frac{\int p(\theta|e_s, b) d\theta}{\int \theta_{e_s} p(\theta|e_s, b) d\theta} = \frac{1}{\mathbb{E}(\Theta_{e_s} | E_s = e_s, B = b)}. \end{aligned}$$

On the other hand, the common approach taken by forensic literature would be to propose the following derivation for the likelihood ratio

$$\begin{aligned} \text{LR} &= \frac{\Pr(E = e | B = b, H = h_p)}{\Pr(E = e | B = b, H = h_d)} = \frac{\Pr(E_c = e_s | B = b, E_s = e_s, H = h_p) \cancel{p(e_s|b, h_p)}}{\Pr(E_c = e_s | B = b, E_s = e_s, H = h_d) \cancel{p(e_s|b, h_d)}} \\ &= \frac{1}{\Pr(E_c = e_s | H = h_d)} = \frac{1}{\theta_{e_s}}. \end{aligned}$$

Then, θ_{e_s} is replaced with $\widehat{\theta_{e_s}} = \mathbb{E}(\theta_{e_s} | B = b)$.

Let us focus on the second to last equality. By looking at the Bayesian network we could already realise that actually E_c is independent only if θ is given. Thus, it is not allowed to simplified E_s in the conditioning unless we are considering frequentist probabilities (\mathcal{Pr}). Last equality is also incorrect. Indeed, it holds that $\Pr(E_c = e_s | B = b, E_s = e_s, h_d) = \int \Pr(E_c = e_s | \Theta = \theta, E_s = e_s, h_d) p(\theta | b, e_s, h_d) d\theta$.

It is true that, in the end, computationally, the difference amounts on using $\mathbb{E}(\theta_{e_s} | B = b, E_s = e_s)$ instead of $\mathbb{E}(\theta_{e_s} | B = b)$ (i.e., the well-known problem of whether to add or

not the suspect to the database before the posterior) and thus the plug-in can be seen as an approximation of the full Bayesian approach. However, it is an hybrid solution, thus conceptually ill defined.

This hybrid approach is often considered Bayesian since the lack of knowledge about θ is dealt with using Bayesian posterior mean $\widehat{\theta}_{e_s} = \mathbf{E}(\Theta_{e_s}|B)$ as a point estimate of θ_{e_s} (Weir, 1996; Curran, 2005; Taroni et al., 2010; Sjerps et al., 2015). This is why we will refer to this way of proceeding as the *Bayesian plug-in method*. As pointed out in Weir (1996), “either the mean or the mode of the posterior distribution can serve as an estimate but each is merely a summary of the whole distribution”. Not only this method is hybrid and inconsistent, but it suffers from several weaknesses. For instance, one would obtain different $\widehat{\mathcal{LR}}$ s depending on whether one wants to estimate θ_{e_s} , $1/\theta_{e_s}$ or $\log_{10}(1/\theta_{e_s})$: this arbitrariness is in some way entailed in the idea of ‘estimating’ the likelihood ratio. Moreover, as stated in Taroni et al. (2015), the likelihood ratio (meaning the Bayesian one) should be calculated, rather than estimated. Including B as part of the data to evaluate, and applying simple Bayesian theory, we can calculate the Bayesian LR, without any estimation needed. Notice that already Foreman et al. (1997) and Brümmer and Swart (2014) proposed a differentiation between the ‘plug-in estimates’ and the ‘full Bayesian analysis’.

4 A useful Lemma

Lemma 4.1 is a result regarding four general random variables A , X , Y , H whose conditional dependencies are represented by the Bayesian network of Figure 5. This is important due to the possibility of applying it to a very common forensic situation: the prosecution and the defence disagree on the the distribution of part of data (Y) but agree on the distribution of the other part (X). The distribution of X and Y depends on some parameter(s) modeled by A .




Figure 5: Conditional dependencies of the random variables of Lemma 4.1

Lemma 4.1. *Given four random variables A , H , X and Y , whose conditional dependencies are represented by the Bayesian network of Figure 5, the likelihood function for h ,*

given $X = x$ and $Y = y$ satisfies

$$\text{lik}(h \mid x, y) \propto \mathbb{E}(p(y \mid x, A, h) \mid X = x).$$

A proof of this lemma can be found in [Cereda \(2015b\)](#). We will see an application of it in Sections 5 and 6.

5 Bayesian LR calculation, based on beta-binomial model

In the binomial model, the database of size N is regarded as the result of a sequence of N Bernoulli trials with parameter θ , where *success* corresponds to the observation of the same type of that observed at the crime scene, and failure to the observation of any other type. Let's denote by b the number of successes among these N experiments. When data is treated as a binomial outcome, the most conventional choice of the prior for the parameter θ (probability of success) is the beta distribution, due to the famous conjugacy property. In forensic and medical statistics literature there are many examples for the use of this distribution for a genetic (autosomal) frequency ([Weir, 1996](#); [Gunel and Wearden, 1995](#); [Roewer et al., 2000](#); [Brenner, 2010](#); [Buckleton et al., 2011](#); [Biedermann et al., 2008, 2013](#)).

$$\Theta \sim \text{Beta}(\alpha, \beta).$$

The observation of the suspect's profile E_s corresponds to another Bernoulli trial, a successful one in the case of interest (the suspect matches the crime stain type). The information provided by the database and the suspect's type can be reduced to the count of profiles of this type in this sample of size $N + 1$ (database and suspect) from the population of interest.

$$B \mid \Theta = \theta \sim \text{Bin}(N, \theta)$$

$$B, E_s \mid \Theta = \theta \sim \text{Bin}(N + 1, \theta)$$

Notice that according to the defence E_c can be seen another Bernoulli experiment of the same kind. On the other hand, according to the prosecution it is equal to 1 with probability one. Stated otherwise,

$$E_c \mid E_s = 1, H = h \sim \begin{cases} \delta_1 & \text{if } H = h_p \\ \theta & \text{if } H = h_d \end{cases}.$$

The Bayesian network of Figure 4 can be used for this model. Hence, we can apply the Lemma 4.1 using $X = (B, E_s)$ (part of data whose distribution is agreed on by defence and distribution) and $Y = E_c$ (part of data whose distribution is disagreed on by defence and distribution). The LR can thus be developed in the following way:

$$\text{LR} = \frac{\mathbb{E}(\Pr(E_c = 1 \mid E_s = 1, H = h_p, \Theta) \mid E_s = 1, B = b)}{\mathbb{E}(\Pr(E_c = 1 \mid E_s = 1, H = h_d, \Theta) \mid E_s = 1, B = b)} = \frac{1}{\mathbb{E}(\Theta \mid E_s = 1, B = b)} = \frac{\alpha + \beta + N + 1}{\alpha + b + 1}. \quad (2)$$

The last equality is due to the fact that, using the well known beta binomial conjugacy property, it holds that

$$\Theta \mid B = b, E_s = 1 \sim \text{Beta}(\alpha + b + 1, \beta + N - b).$$

The LR as in (2), also proposed in Dawid and Mortera (1996) and Taroni et al. (2015), can be compared to the one obtained with the ‘standard’ Bayesian plug-in estimate (Weir, 1996; Taroni et al., 2010):

$$\widehat{\text{LR}} = \frac{\alpha + \beta + N}{\alpha + b}.$$

It is easy to see that the Bayesian plug-in $\widehat{\text{LR}}$ is a non conservative estimate of LR, in a way that is unfavourable to the defence. Indeed, $\text{LR} < \widehat{\text{LR}} \Leftrightarrow \beta + N > b$, which is always true, since $b \leq N$ and $\beta > 0$. Notice that there is an alternative derivation for (2). It can be obtained in a two step evaluation: first, the observation of the database B and of the suspect haplotype E_s updates the probability Pr , then the updated probability is used to calculate the LR for the observation of another identical haplotype (the one found at the crime scene).

First step The probability Pr is updated to $\text{Pr}^{**}(\cdot) = \text{Pr}(\cdot \mid E_s = 1, B = b)$ after the database and the haplotype of the suspect are observed. In practice, the prior distribution $\text{Beta}(\alpha, \beta)$ on θ is updated to the posterior $\text{Beta}(\alpha + 1, \beta + N - 1)$.

Second step The new probability Pr^{**} is used to calculate the likelihood ratio for the observation of the haplotype from the crime scene:

$$\text{LR} = \frac{\text{Pr}^{**}(E_c = 1 \mid H = h_p)}{\text{Pr}^{**}(E_c = 1 \mid H = h_d)} = \frac{1}{\mathbb{E}^{**}(\Theta)} = \frac{1}{\frac{\alpha + b + 1}{\alpha + \beta + 1 + N}}.$$

Sensitivity analysis. The sensitivity of the quantities $\log_{10} \text{LR}$, $\log_{10} \widehat{\text{LR}}$, and of the difference between them, to the hyperparameters α and β of the beta prior is shown in Figure 6, for the rare type case (i.e., $b = 0$), and with $N = 100$. In particular, the figure shows the variation of $\log_{10} \text{LR}$ (a), of the plug-in estimate $\log_{10} \widehat{\text{LR}} = \log_{10} \widehat{\text{LR}}$ (b), and of the difference $\log_{10} \widehat{\text{LR}} - \log_{10} \text{LR}$ (c), when different values of α (x axis) and β (only five values corresponding to the different lines) are chosen in the interval $(0, 20]$.

Observing Figure 6 (or analysing (2)), it can be seen that the three quantities of interest hardly depend on β , while they decrease as α increases. In particular, when α decreases to 0, $\log_{10} \text{LR}$ behaves as $\log_{10}(1 + \beta + N)$, while $\log_{10} \widehat{\text{LR}}$ increases to $+\infty$. Another way to see this is that, for fixed β , as α increases, the prior distribution on θ resembles more and more to the degenerate distribution localised on the value $\theta = 1$ (notice this is inappropriate for the rare type match case). This means that the haplotype whose




Figure 6: Sensitivity analysis for the three quantities $\log_{10} \text{LR}$ (a), $\log_{10} \widehat{\text{LR}}$ (b), $\log_{10} \widehat{\text{LR}} - \log_{10} \text{LR}$ (c), in the beta binomial model, when $\alpha \in (0, 20]$ (x axis), and for $\beta \in \{1, 5, 10, 15, 20\}$ (corresponding to the different lines, where highest line corresponds to highest β).

population proportion is modelled through the random variable Θ (i.e., the haplotype of the crime stain and of the suspect) has probability one to be observed, which leads to $\widehat{\text{LR}} = 1$ (hence, $\log_{10} \text{LR} = 0$). On the other hand, if α decreases to zero, the prior distribution over θ tends to resemble to the degenerate distribution localised on the value $\theta = 0$. This leads to $\widehat{\text{LR}} = 1/0 = +\infty$. On the whole, the plug-in estimate $\widehat{\text{LR}}$ is less stable than LR, as can be seen comparing Figures 6 (a) and (b), in the sense that is more sensitive to changes in α (especially for small values). The difference, represented in (c) has, for fixed β , a vertical asymptote when $\alpha \rightarrow 0$, increasing as fast as $\log_{10} 1/\alpha$. On the other hand it decreases to 0 with an horizontal asymptote when $\alpha \rightarrow \infty$. From Figure 6 (c) it can be observed that the difference is important only for small values of α . Otherwise the two methods would lead essentially to the same conclusions, so that the plug-in can be seen as a good approximation of the proper Bayesian procedure.

6 Bayesian LR calculation, based on Dirichlet-multinomial model

When database is treated as a multinomial sample of size N from a population with k different haplotypes, the conventional choice of the prior for the vector θ containing the population frequencies of all the different haplotypes in Nature is the Dirichlet distribution. Literature provides many examples of the use of this method for the frequencies of autosomal markers (Curran et al., 2002; Balding, 1995; Lange, 1995; Weir, 1996; Buckleton and Curran, 2005; Taroni et al., 2010, e.g.). However, these methods don't consider the uncertainty about the number k of possible types in the population, and this can be a problem especially since we want to apply it to Y-STR haplotypes, for which the database

often does not offer a good coverage. If, in addition, we are using the model for the rare type match case, then we have to find a solution. The problem of estimating k is a very challenging one. It has been addressed both with frequentist methods (Chao and Lee, 1992; Haas and Stokes, 1998, e.g.) and with Bayesian methods (Hill, 1968; Lewins and Joanes, 1984; Barger and Bunge, 2010, e.g.). We propose the derivation of a full Bayesian LR which uses priors over the number k of different types in the population. The model is represented by the Bayesian network of Figure 7. The bottom part (from node Θ down) has a well-known structure (see Figure 4), while the upper part needs further explanation.




Figure 7: Bayesian network for Dirichlet-multinomial model, when k is randomized.

Assume that there may be at most m theoretically possible profiles alphabetically² ordered in a vector, called \mathbf{s} . For instance, $m = 20^{10}$ (10 loci, with 20 possible alleles each). Only k of them are actually present in Nature (or more specifically in the population of interest), but k is not known and also which of the m are those k is not known.

We will denote as K the random variable which represents how many of the m potentially possible haplotypes are actually present in the population of interest. The prior distribution for k is denoted generically as $p(k)$. The random vector **Type**, of length k , contains the ordered positions, in vector \mathbf{s} , of the k haplotypes of the population of interest. A particular configuration of **Type** is denoted as $\mathbf{t} = (i_1, \dots, i_k)$, where $i_1 < \dots < i_k$. \mathbf{t} is chosen uniformly at random from the possible $\binom{m}{k}$ configurations. The random vector Θ contains the population proportions of all the haplotypes, both those whose position is contained in **Type**, and those that are not (corresponding to zero entries). A particular configuration of Θ is denoted $\theta = (\theta_1, \dots, \theta_m)$, many entries of which are zero. We assume that the positive entries, i.e., $(\theta_i \mid i \in \mathbf{t})$, are drawn from a k dimensional Dirichlet

²Remember each profile is a list of numbers.

distribution with all k hyper parameters α equal to 1. Now, as usual, H represents the hypotheses of interest, and can take the value $h \in \{h_p, h_d\}$, according to the prosecution or the defence, respectively. E_s and E_c contain the index e_s and e_c of the haplotypes of the suspect and of the crime scene, respectively. In the situation of interest $e_c = e_s$. Lastly, random vector \mathbf{B} represents the database, seen as a multinomial sample from the population with parameters N and $\boldsymbol{\theta}$. A particular configuration of \mathbf{B} is denoted $\mathbf{b} = (b_1, \dots, b_m)$ representing the absolute frequency in the database of each of the m haplotypes. It contains $k_{obs} < k$ positive values, and many zeros. By applying Lemma 4.1 to this situation we have that

$$\text{LR} = \frac{\mathbb{E}(\Pr(E_c = e_s \mid E_s = e_s, \mathbf{B} = \mathbf{b}, \boldsymbol{\Theta}, H = h_p) \mid E_s = e_s, \mathbf{B} = \mathbf{b})}{\mathbb{E}(\Pr(E_c = e_s \mid E_s = e_s, \mathbf{B} = \mathbf{b}, \boldsymbol{\Theta}, H = h_d) \mid E_s = e_s, \mathbf{B} = \mathbf{b})} = \frac{1}{\mathbb{E}(\Theta_{e_s} \mid E_s = e_s, \mathbf{B} = \mathbf{b})}.$$

It can be proven that for $\alpha = 1$ this leads to

$$\text{LR} = \frac{\frac{1}{2} \sum_{k=k_{obs}+1}^m \binom{k}{k_{obs}+1} \frac{\Gamma(k)}{\Gamma(k+N+1)} p(k)}{\sum_{k=k_{obs}+1}^m \binom{k}{k_{obs}+1} \frac{\Gamma(k)}{\Gamma(k+N+2)} p(k)} \quad (3)$$

Notice that the likelihood ratio depends on the data only through k_{obs} . This is due to the choice of the symmetric Dirichlet prior, and of the uniform prior for **Type**. In particular, this tells us that data can be reduced by sufficiency to k_{obs} . The likelihood ratio obtained through a classical plug-in Bayesian estimation is:

$$\widehat{\text{LR}} = \frac{\bar{k}\alpha + N}{\alpha + b_{e_s}} = \bar{k} + N. \quad (4)$$

where the number of haplotypes is a fixed value \bar{k} , to be chosen (or estimated) in advance. In order to compare the two values (4) and (3) we need to choose a value for \bar{k} . A reasonable choice can be $\bar{k} = \mathbf{E}(K)$. Among the possible priors one can choose for K , we decided to test the Poisson distribution (see Section 6.1) and the negative binomial distribution (see Section 6.2).

6.1 Poisson prior

In this section a Poisson distribution with parameter λ , truncated so as to have support only on $\{1, 2, \dots, m\}$, is chosen as prior distribution for K .

$$p(k) := p(k; \lambda) \propto \begin{cases} \frac{e^{-\lambda} \lambda^k}{k!} & \text{if } k \in \{1, \dots, m\} \\ 0 & \text{elsewhere} \end{cases}$$

where $\lambda > 0$. If λ and m are large enough, the normalising constant can be omitted and we have the standard Poisson distribution:

$$p(k; \lambda) = \frac{e^{-\lambda} \lambda^k}{k!}, \quad \forall k \in \mathbb{N}$$

The LR in (3) becomes

$$\text{LR} = \frac{1}{2} \frac{\sum_{k=k_{obs}+1}^m \frac{\lambda^k}{k-k_{obs}-1!} \frac{\Gamma(k)}{\Gamma(N+k+1)}}{\sum_{k=k_{obs}+1}^m \frac{\lambda^k}{k-k_{obs}-1!} \frac{\Gamma(k)}{\Gamma(N+k+2)}}$$

It is then of interest to analyse the quantities $\log_{10} \text{LR}$, $\log_{10} \widehat{\text{LR}}$, and the difference $\log_{10} \widehat{\text{LR}} - \log_{10} \text{LR}$ between them and to carry on a sensitivity analysis to see how these quantities vary when parameter λ changes.

Sensitivity analysis In the rare type match problem (i.e., $b_{es} = 0$), when a $\text{Poisson}(\lambda)$ prior is chosen for the dimension K of the Dirichlet distribution (with all parameters α equal to 1), the sensitivity of the three quantities $\log_{10} \text{LR}$, $\log_{10} \widehat{\text{LR}}$, and of their difference, to λ and k_{obs} is shown in Figure 8, when $N = 100$.




Figure 8: Poisson prior, for the Dirichlet model. Sensitivity analysis of $\log_{10} \text{LR}$ (a, black lines), $\log_{10} \widehat{\text{LR}}$ (a, dashed line), and of the difference $\log_{10} \widehat{\text{LR}} - \log_{10} \text{LR}$ (b), to different values of $\lambda \in [1, 10\,000]$ (x axis), and $k_{obs} \in \{30, 40, 50, 60, 70, 80, 90, 100\}$ (represented by the different lines where highest line corresponds to highest β).

In particular, it can be inferred that the LR depends little on k_{obs} and a lot on λ . When λ is big (which is typically true λ being the expected value of the number of different Y-STR haplotypes in a population) the LR depends almost only on λ . In particular, LR increases linearly with λ , since $\text{LR} \sim \lambda/2$. This can be explained by replacing the Poisson prior on k , by the degenerate distribution localised on (the integer part of) λ : $f_K(k) = f(k; \lambda) = \mathbb{1}_{\{\lambda\}}(k)$, for $\lambda \in \{1, 2, \dots\}$. This approximation is sensible for large values of λ in virtue of the law of large numbers (the $\text{Poisson}(\lambda)$ being the sum of λ

Poisson(1)). In this case (3) becomes

$$\text{LR} = \frac{1 + N + \lambda}{2} \sim \frac{\lambda}{2}, \text{ for } \lambda \rightarrow +\infty, \text{ and } N \text{ fix.}$$

The plug-in estimates of $\log_{10} \widehat{\text{LR}}$ (as defined in (4) and with the choice of $\bar{k} = \lambda$) is the dashed line shown in Figure 8 (a). The difference between the ‘true’ value $\log_{10} \text{LR}$, and the estimated one $\log_{10} \widehat{\text{LR}}$ is shown in Figures 8 (b). In particular, one can see that, for big λ , it decreases when λ increases and depends a little on k_{obs} , while for small values of λ it has the opposite behaviour, and depends more strongly on k_{obs} . Note that, again, the plug-in method overestimates the LR by up to almost half of an order of magnitude.

6.2 Negative binomial prior

A different choice is that of using as prior for k the negative binomial distribution (Hill, 1968, 1979; Lewins and Joanes, 1984) as prior distribution for K . For our model a negative binomial distribution truncated so as to have support $\{1, \dots, m\}$ is more appropriate. It is defined as:

$$\Pr(K = k|r, q) \propto \begin{cases} \binom{k+r-1}{k} (1-q)^k q^r & \text{if } k \in \{1, \dots, m\} \\ 0 & \text{elsewhere} \end{cases}$$

where $r > 0$ and $q \in (0, 1)$. However, if $\mathbb{E}(K)$ is large, but small compared to m , the normalise factor is almost 1 and the standard negative binomial distribution can be used as prior distribution over K :

$$\Pr(K = k|r, q) = \binom{k+r-1}{k} (1-q)^k q^r, \quad \forall k \in \mathbb{N}.$$

Using this prior, the likelihood ratio in (3) becomes:

$$\text{LR} = \frac{1 \sum_{k=k_{obs}+1}^m (1-q)^k \frac{\Gamma(k)}{\Gamma(k+N+1)} \frac{\Gamma(k+r)}{\Gamma(k-k_{obs})}}{2 \sum_{k=k_{obs}+1}^m (1-q)^k \frac{\Gamma(k)}{\Gamma(k+N+2)} \frac{\Gamma(k+r)}{\Gamma(k-k_{obs})}}. \quad (5)$$

In the following, a series of properties of the (zero truncated) negative binomial distribution will be listed, which will help to understand why this choice is more appropriate than the choice of the Poisson distribution as a prior for K . We will denote as $\text{NB}(r, q)$ a random variable distributed according to a negative binomial with parameters r and q , and $\text{P}(\lambda)$ a random variable distributed according to a Poisson distribution with parameter λ .

1. The mean and variance of $\text{NB}(r, q)$ are, respectively, $\mathbf{E}(\text{NB}(r, q)) = (1-q)r/q$ and $\mathbf{Var}(\text{NB}(r, q)) = (1-q)r/q^2$. This represents an advantage over the use of a Poisson distribution where these two quantities can’t be tuned independently one another, since $\mathbf{E}(\text{P}(\lambda)) = \mathbf{Var}(\text{P}(\lambda)) = \lambda$. Thus, the use of a negative binomial prior guarantees more flexibility.
2. The negative binomial $\text{NB}(r, q)$ is a Gamma mixture of Poisson.

3. For fixed $\lambda = \mathbf{E}(\text{NB}(r, q))$, when r increases, the negative binomial $\text{NB}(r, q)$ tends in distribution to $P(\lambda)$. This means that the negative binomial distribution can be seen as an extension of the Poisson distribution.

The same properties apply to the $[0, m]$ -truncated case, both for the Negative Binomial, and for the Poisson, if m is big enough **and the probability of 0 is small**.

6.3 Sensitivity analysis

A classical approach to sensitivity analysis for the negative binomial prior would be to analyse the sensitivity of $\log_{10} \text{LR}$ to changes of r and q , and k_{obs} , the three parameters appearing in (5). However, we decided to use as parameters r , k_{obs} (the number of different haplotypes observed in the database) and λ , the mean value of the negative binomial prior. In this way it is easier to see how the results depend on the average number of haplotypes in Nature, and that for big r we fall back in the Poisson case, as explained in property 3. Figure 1 represents the sensitivity analysis for $\log_{10} \text{LR}$, $\log_{10} \widehat{\text{LR}}$ and the difference $\log_{10} \widehat{\text{LR}} - \log_{10} \text{LR}$ in the rare type case ($b_{i_s} = 0$) for $\alpha = 1$, $N = 100$.

It can be inferred from this analysis that when r increases the values depend more and more on λ and less and less on k_{obs} .

According to the second column of Figure 1, one can see that for $r \geq 100$ the plug-in estimate always exceeds $\log_{10} \text{LR}$. Anyway, the difference is only significant if r is small, in particular for high values of λ .

6.4 Remarks about conventional priors

As mentioned above, the beta distribution and the Dirichlet distribution are the conventional choices for the prior in case of binomial or multinomial model, respectively. As stated in Curran et al. (2002), this ‘remains the accepted standard in some laboratories’ because of the ‘appeal of simplicity and ease of implementation’. Although we agree that this may have been a very sensible reason some decades ago, nowadays, with the computational skill provided by our computers, there are no more excuses to limit ourselves to these convenient priors. Indeed, a prior should reflect the expert beliefs rather than standards of computational ease.

For the beta prior, the dependencies of the LR results on the value of the hyper parameter α stresses once more the need of a different choice. Moreover, the model is profile-specific, meaning that the beta priors is supposed to model the frequency of the profile observed at the crime scene. The model is thus to be defined after the data have been observed, and this seems to contradict common Bayesian principles.

For the Dirichlet prior, we have a similar issue. The dimension k of this prior should correspond to the number of different DNA types in the population. This problem, which

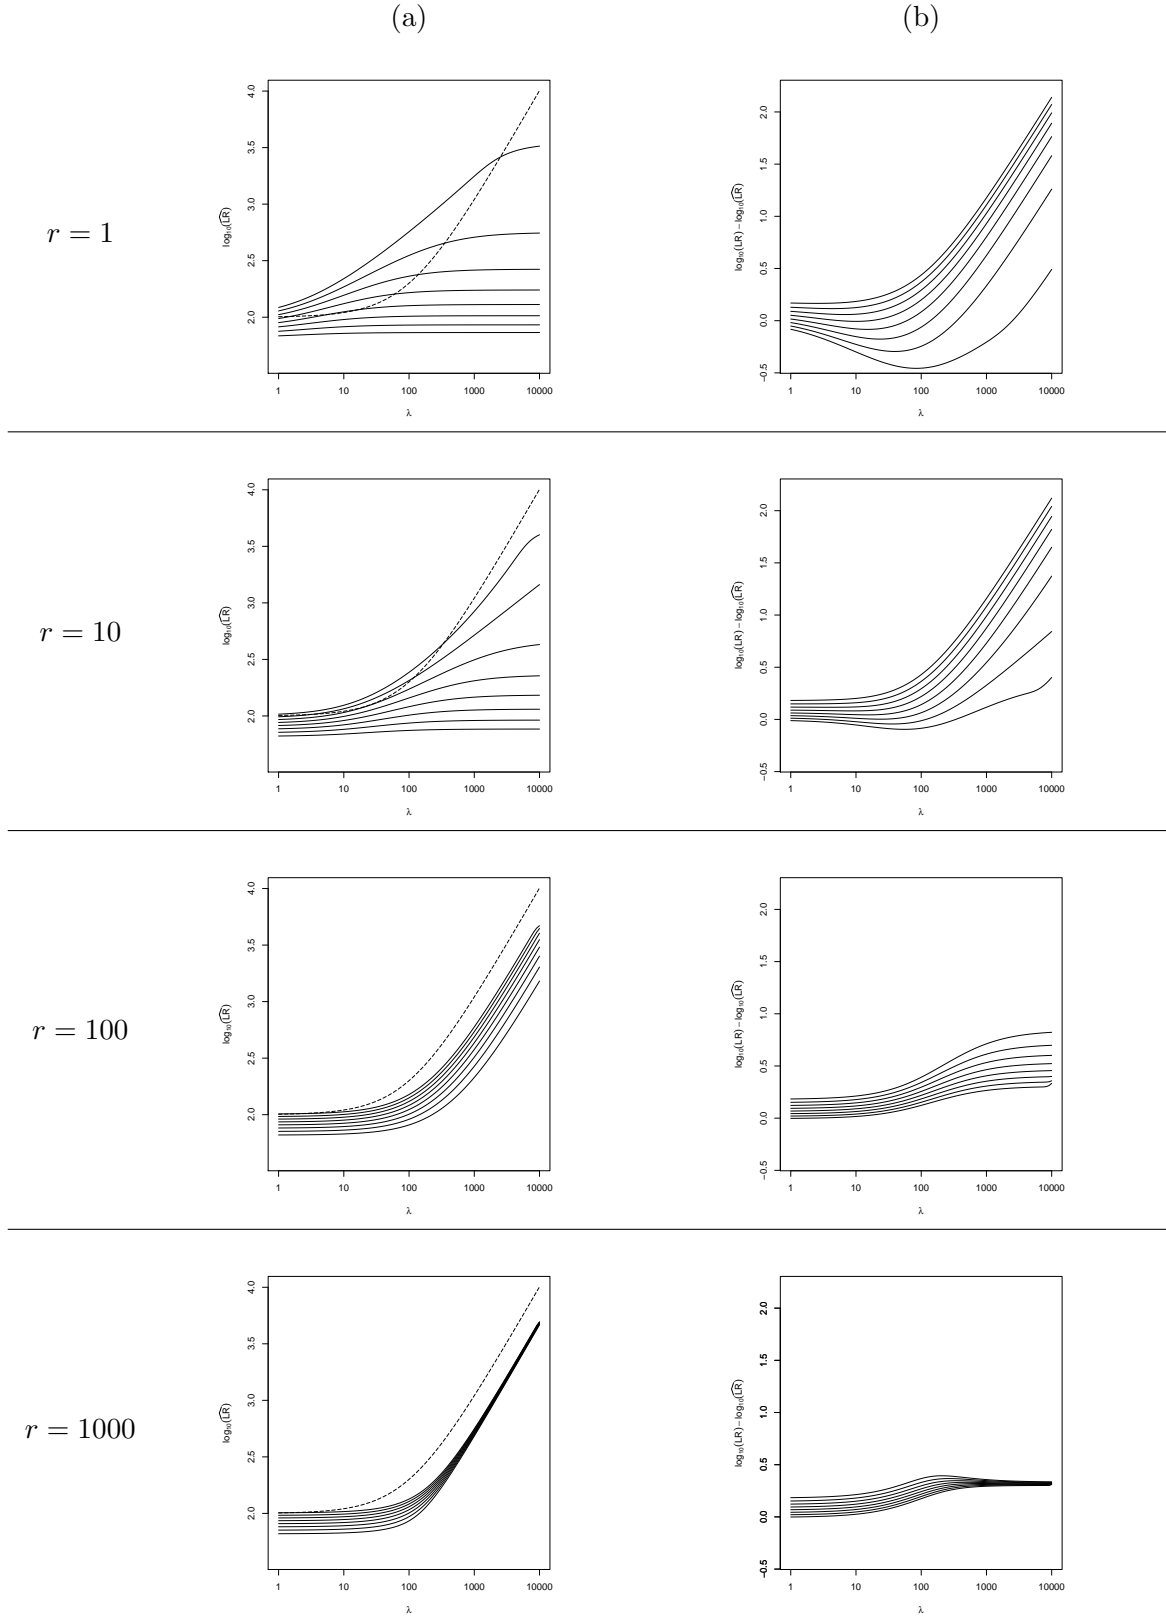


Table 1: Sensitivity analysis for the three quantities $\log_{10} \text{LR}$ (first column, black lines), $\log_{10} \widehat{\text{LR}}$ (first column, dashed line) and the difference $\log_{10} \widehat{\text{LR}} - \log_{10} \text{LR}$ (second column) to different values of $\lambda = \mathbf{E}(K)$ (x axis) and of $k_{obs} \in \{30, 40, 50, 60, 70, 80, 90, 100\}$ (represented by the different lines, where highest line corresponds to the highest k_{obs}).

for autosomal markers could be easily overcome, is more important when Y-STR haplotypes are considered, the state space being huge, and the database hardly representative. If we choose as k the number of different types observed in the database, then we are in trouble every time a new haplotype is observed, as for the situation of interest for this paper. By treating k as a Bayesian would do for an unknown quantity, we expected the likelihood ratio to depend a lot on the mean value of the prior chosen for k . The Dirichlet method with all parameters $\alpha = 1$, and with a prior over K , turned out to depend only on the number of observed haplotypes in the database (and not on their frequencies). This is actually unattractive for Y-STR data, and is due to the symmetry. The data does not overrule the prior which makes *all* the positive p_i the same in size, and it is also the reason why the likelihood ratios obtained using the two methods (beta-binomial, and Dirichlet-multinomial) do not differ too much. Notice that for this prior we only focused on the case in which all the parameters α are equal to 1. More could have been done, for instance explore the sensitivity of the likelihood ratio to changes in the α (Triggs and Curran, 2006), or use hierarchical model (Chen and McGee, 2008). However, we preferred to investigate other types of prior (Cereda, 2015c) which we believe are more appropriate for Y-STR haplotypes frequencies.

The two methods of Section 5 and Section 6 differ in the choice of information retained from the database. The Beta method only retains as information the frequency of the observed haplotype. A lot of information regarding other haplotypes is discarded, such as how many have been observed, and their frequencies. Let us point out that if there will ever be guidelines on how to choose the hyper parameters of the beta prior and of the Dirichlet prior, they should be compatible: hence the beta prior should be the one obtained from the Dirichlet by marginalisation.

7 Conclusion

This paper is intended to have several take-home messages. The first one is that a forensic statistician before starting any evaluation should make up his mind if he wants to use frequentist or Bayesian methods, since we have seen that the corresponding likelihood ratios are differently defined. If a Bayesian approach is chosen, which has the advantage that everything is combined into a single number, without any uncertainty involved, the LR should be calculated in a principled way. Bayesian plug-in (and frequentist plug-in), often proposed as proper Bayesian approach, can sometimes be seen as a convenient approximation of the Bayesian LR, but this paper has shown that the full Bayesian method is not more difficult. Moreover, the Bayesian plug-in is almost always anti-conservative in a way that is unfair to defence, and there are sometimes significant differences with the full Bayesian method for particular choices of the hyper-parameters of the priors.

All this has been shown when the conventional choices for the priori (beta or Dirichlet) are made. The choice of the prior is an issue indeed. We believe that a true Bayesian should not make use of conventional priors, but of his own priors. Especially because, as shown, this conventional choice leads to likelihood ratios which strongly depend on the hyperparameters of these priors. Choosing more realistic prior may increase the difficulty of the computation of the likelihood ratio, but, also thanks to modern computational tools, this should not stop people from preferring them.

Acknowledgements

I am indebted to Charles Brenner for the useful discussion about this paper, which lead to many improvements. This research was supported by the Swiss National Science Foundation, through grants no. 105311-1445570 and 10531A-156146/1, and carried out in the context of a joint research project, supervised by Franco Taroni (University of Lausanne, Ecole des sciences criminelles), and Richard Gill (Mathematical Institute, Leiden University).

References

- Aitken, C. and Taroni, F. (2004). *Statistics and the Evaluation of Evidence for Forensics Scientists*. John Wiley & Sons, Chichester.
- Andersen, M. M., Eriksen, P. S., and Morling, N. (2013). The discrete Laplace exponential family and estimation of Y-STR haplotype frequencies. *Journal of Theoretical Biology*, 329:39–51.
- Anevski, D., Gill, R. D., and Zohren, S. (2013). Estimating a probability mass function with unknown labels. <http://arxiv.org/abs/1312.1200>.
- Balding, D. (2005). *Weight-of-evidence for Forensic DNA Profiles*. John Wiley & Sons Hoboken, NJ.
- Balding, D. J. (1995). Estimating products in forensic identification using DNA profiles. *Journal of the American Statistical Association*, 90:839–844.
- Barger, K. and Bunge, J. (2010). Objective Bayesian estimation for the number of species. *Bayesian Analysis*, 5:765–785.
- Biedermann, A., Garbolino, P., and Taroni, F. (2013). The subjectivist interpretation of probability and the problem of individualisation in forensic science. *Science & Justice*, 53:192–200.

- Biedermann, A., Taroni, F., Bozza, S., and Aitken, C. (2008). Analysis of sampling issues using Bayesian networks. *Law, Probability and Risk*, 7:35–60.
- Brenner, C. H. (2010). Fundamental problem of forensic mathematics—The evidential value of a rare haplotype. *Forensic Science International: Genetics*, 4:281–291.
- Brümmer, N. and Swart, A. (2014). Bayesian calibration for forensic evidence reporting. arXiv:1403.5997.
- Buckleton, J. and Curran, J. (2005). Sampling effects. In Buckleton, J., Triggs, C., and Walsh, S. J., editors, *Forensic DNA evidence interpretation*, chapter 6, pages 197–216. CRC Press, Boca Raton.
- Buckleton, J., Krawczak, M., and Weir, B. (2011). The interpretation of lineage markers in forensic DNA testing. *Forensic Science International: Genetics*, 5:78 – 83.
- Cereda, G. (2015a). Impact of model choice on LR assessment in case of rare haplotype match (bayesian approach). arXiv:1502.02406.
- Cereda, G. (2015b). Impact of model choice on LR assessment in case of rare haplotype match (frequentist approach). arXiv:1502.04083.
- Cereda, G. (2015c). Non parametric bayesian approach to LR assessment in case of rare haplotype match. arxiv:1506.08444.
- Cereda, G., Biedermann, A., Hall, D., and Taroni, F. (2014). An investigation of the potential of DIP-STR markers for DNA mixture analyses. *Forensic Science International: Genetics*, 11:229–240.
- Chao, A. and Lee, S.-M. (1992). Estimating the number of classes via sample coverage. *Journal of the American Statistical Association*, 87:210–217.
- Chen, Z. and McGee, M. (2008). A bayesian approach to zero numerator problems using hierarchical models. *Journal of Data Science*, 6:261–268.
- Curran, J., Buckleton, J., Triggs, C., and Weir, B. (2002). Assessing uncertainty in DNA evidence caused by sampling effects. *Science & Justice*, 42:29–37.
- Curran, J. M. (2005). An introduction to Bayesian credible intervals for sampling error in DNA profiles. *Law, Probability and Risk*, 4:115–126.
- Dawid, A. P. and Mortera, J. (1996). Coherent analysis of forensic identification evidence. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58:425–443.

- Egeland, T. and Salas, A. (2008). Estimating haplotype frequency and coverage of databases. *PLoS ONE*, 3:e3988–e3988.
- Evett, I. and Weir, B. (1998). *Interpreting DNA evidence: Statistical Genetics for Forensic Scientists*. Sinauer Associates, Sunderland.
- Favaro, S., Lijoi, A., Mena, R. H., and Pruenster, I. (2009). Bayesian nonparametric inference for species variety with a two parameter poisson-dirichlet process prior. *Journal of the Royal Statistical Society: Series B (Methodological)*, 71:993–1008.
- Foreman, L., Smith, A., and Evett, I. (1997). Bayesian analysis of DNA profiling data in forensic identification applications. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 160:429–459.
- Gale, W. A. and Church, K. W. (1994). What’s wrong with adding one? In *Corpus-Based Research into Language*. Rodolpi.
- Gill, P., Brenner, C., Brinkmann, B., Budowle, B., Carracedo, A., Jobling, M. A., de Knijff, P., Kayser, M., Krawczak, M., Mayr, W. R., Morling, N., Olaisen, B., Pascali, V., Prinz, M., Roewer, L., Schneider, P. M., Sajantila, A., and Tyler-Smith, C. (2001). DNA Commission of the International Society of Forensic Genetics: recommendations on forensic analysis using Y-chromosome STRs. *Forensic Science International*, 124:5–10.
- Good, I. (1953). The population frequencies of species and the estimation of population parameters. *Biometrika*, 40:237–264.
- Gunel, E. and Wearden, S. (1995). Bayesian estimation and testing of gene frequencies. *Theoretical and Applied Genetics*, 91:534–543.
- Haas, P. J. and Stokes, L. (1998). Estimating the number of classes in a finite population. *Journal of the American Statistical Association*, 93:1475–1487.
- Hill, B. M. (1968). Posterior distribution of percentiles: Bayes’ theorem for sampling from a population. *Journal of the American Statistical Association*, 63:677–691.
- Hill, B. M. (1979). Posterior moments of the number of species in a finite population and the posterior probability of finding a new species. *Journal of the American Statistical Association*, 74:668–673.
- van der Hout, A. and Alberink, I. (2015). Posterior distributions for likelihood ratios in forensic science. <http://www.ucl.ac.uk/ucaakadl/LR2015.pdf>
- Jovanovic, B. D. and Levy, P. S. (1997). A look at the rule of three. *The American Statistician*, 51:137–139.

- Krichevsky, R. and Trofimov, V. (1981). The performance of universal coding. *IEEE Transactions on Information Theory*, 27:199–207.
- Lange, K. (1995). Applications of the Dirichlet distribution to forensic match probabilities. *Genetica*, 96:107–117.
- Laplace, P. (1814). *Essai philosophique sur les probabilités*. m.me v.e Courcier.
- Lewins, W. A. and Joanes, D. N. (1984). Bayesian estimation of the number of species. *Biometrics*, 40:323–328.
- Lijoi, A., Mena, R. H., and Pruenster, I. (2007). Bayesian nonparametric estimation of the probability of discovering new species. *Biometrika*, 94:769–786.
- Louis, T. A. (1981). Confidence intervals for a binomial parameter after observing no successes. *The American Statistician*, 35(3):154–154.
- Lucy, D. (2005). *Introduction to Statistics for Forensic Scientists*. John Wiley & Sons, Chichester.
- Orlitsky, A., Santhanam, N., and Zhang, J. (2003). Always Good Turing: asymptotically optimal probability estimation. *Science (New York, N.Y.)*, 302:427–431.
- Orlitsky, A., Santhanam, N. P., Viswanathan, K., and Zhang, J. (2004). On Modeling Profiles Instead of Values. In *UAI*, pages 426–435.
- Roewer, L., Kayser, M., de Knijff, P., Anslinger, K., Betz, A., Caglia, A., Corach, D., Furedi, S., Henke, L., Hidding, M., Kargel, H., Lessig, R., Nagy, M., Pascali, V., Parson, W., Rolf, B., Schmitt, C., Szibor, R., Teifel-Greding, J., and Krawczak, M. (2000). A new method for the evaluation of matches in non-recombining genomes: application to Y-chromosomal short tandem repeat (STR) haplotypes in European males. *Forensic Science International*, 114:31–43.
- Sjerps, M. J., Alberink, I., Bolck, A., Stoel, R. D., Vergeer, P., and van Zanten, J. H. (2015). Uncertainty and Ir: to integrate or not to integrate, that’s the question. *Law, Probability and Risk*.
- Taroni, F., Biedermann, A., Bozza, S., Garbolino, P., and Aitken, C. (2014). *Bayesian networks for probabilistic inference and decision analysis in forensic science*. John Wiley & Sons, Chichester, second edition edition.
- Taroni, F., Bozza, S., Biedermann, A., and Aitken, C. (2015). Dismissal of the illusion of uncertainty in the assessment of a likelihood ratio. In press.

- Taroni, F., Bozza, S., Biedermann, A., Garbolino, P., and Aitken, C. (2010). *Data Analysis in Forensic Science: A Bayesian Decision Perspective*. Statistics in Practice. Wiley, Chichester.
- Tiwari, R. C. and Tripathi, R. C. (1989). Nonparametric bayes estimation of the probability of discovering a new species. *Communications in statistics: Theory and methods*, A18:877–895.
- Triggs, C. M. and Curran, J. M. (2006). The sensitivity of the Bayesian hpd method to the choice of prior. *Science & Justice*, 46:169–178.
- Weir, B. (1996). *Genetic data analysis 2*. Sinauer Associates, Sunderland.
- Winkler, R. L., Smith, J. E., and Fryback, G. (2002). The role of informative priors in zero-numerator problems: being conservative versus being candid. *The American Statistician*, 64:1–4.